

On Differential Photometric Reconstruction for Unknown, Isotropic BRDFs

Manmohan Chandraker

Jiamin Bai

Ravi Ramamoorthi

Abstract—This paper presents a comprehensive theory of photometric surface reconstruction from image derivatives, in the presence of a general, unknown isotropic BRDF. We derive precise topological classes up to which the surface may be determined and specify exact priors for a full geometric reconstruction. These results are the culmination of a series of fundamental observations. First, we exploit the linearity of chain rule differentiation to discover photometric invariants that relate image derivatives to the surface geometry, regardless of the form of isotropic BRDF. For the problem of shape from shading, we show that a reconstruction may be performed up to isocontours of constant magnitude of the gradient. For the problem of photometric stereo, we show that just two measurements of spatial and temporal image derivatives, from unknown light directions on a circle, suffice to recover surface information from the photometric invariant. Surprisingly, the form of the invariant bears a striking resemblance to optical flow, however, it does not suffer from the aperture problem. This photometric flow is shown to determine the surface up to isocontours of constant magnitude of the surface gradient, as well as isocontours of constant depth. Further, we prove that specification of the surface normal at a single point completely determines the surface depth from these isocontours. In addition, we propose practical algorithms that require additional initial or boundary information, but recover depth from lower order derivatives. Our theoretical results are illustrated with several examples on synthetic and real data.



1 INTRODUCTION

The image formation process is an interplay between the geometry of a scene, its reflectance properties and the illumination conditions under which it is observed. Photometric reconstruction methods seek to recover shape from images using shading cues, possibly under varying illumination. For Lambertian scenes, there exist wide-ranging theories of photometric stereo and shape-from-shading to recover shape from image intensities and gradients, respectively. However, the theoretical relationships between image gradients and surface geometry, for general BRDFs and unknown light source motions, remains largely unexplored.

This paper presents a fundamental study of the utility of image gradients for scene inference in the presence of complex BRDFs and unknown light directions. We delineate exact topological classes up to which scene structure may be determined when BRDF and lighting are unknown. We also specify simple geometric priors that suffice for further disambiguating the full Euclidean structure and present practical reconstruction algorithms for the same. The various aspects of our theory are illustrated in Figure 1 for the problem of photometric stereo.

We begin with the observation in Section 4 that a single image formation equation yields several differential equations when spatial and temporal derivatives of the image are considered. Moreover, these differential equations are linear in the BRDF derivatives, so under appropriate conditions (such as circular motion of the light source), they may be eliminated

to yield a photometric invariant that relates image derivatives, surface geometry and light source directions.

As an initial demonstration, we show in Section 5 that for the problem of shape from shading with a collocated light source, such an invariant is easy to derive. The form of the invariant readily indicates that for an unknown BRDF, surface reconstruction is possible, at best, up to the level curves of the magnitude of the gradient. The major theoretical development in the paper, covering Sections 6 to 8, focuses on the problem of photometric stereo, where we assume that a differential pair of lights undergo unknown motion on a circle. In a series of results, we show that the seemingly complex form of the corresponding invariant, upon careful analysis, yields a rich description of the surface structure.

A surprising discovery in Section 6 is that for isotropic BRDFs, it is possible to uncover such an invariant that is independent of lighting directions. Thus, we do not require calibrated measurements of the circular motion of the differential pair. Moreover, we show that the $N \times 3$ matrix whose columns are the image derivatives at a pixel under N light source positions, must be rank 2 and its null vector, $(\lambda, \kappa, 1)^\top$, is determined by surface geometry alone. The form of this invariant bears a striking resemblance to the optical flow relation, so we call it *photometric flow*. However, in contrast to optical flow, photometric flow stems from correctly accounting for and eliminating complex BRDF effects, rather than any physically inaccurate brightness constancy assumptions.

For a surface $z(x, y) : \mathbb{R}^2 \rightarrow \mathbb{R}$, the entities λ and κ are space varying functions whose relationship to surface depths, z , or gradients ∇z , is not straightforward at first glance. Yet, we show in Section 7 that the information contained in those functions can be succinctly expressed as the direction of the surface gradient and the direction of the gradient of the scalar field $\|\nabla z\|$. Thus, from image information alone, one may determine surface structure up to level curves of constant depth

• M. Chandraker is with NEC Laboratories America, at Cupertino, CA 95014. E-mail: manu@sv.nec-labs.com

• J. Bai and R. Ramamoorthi are with the EECS Department at the University of California, Berkeley, at Berkeley, CA 94720. Email: {bjiamin, ravir}@eecs.berkeley.edu

This work was performed when the first author was a postdoctoral researcher at the University of California, Berkeley.

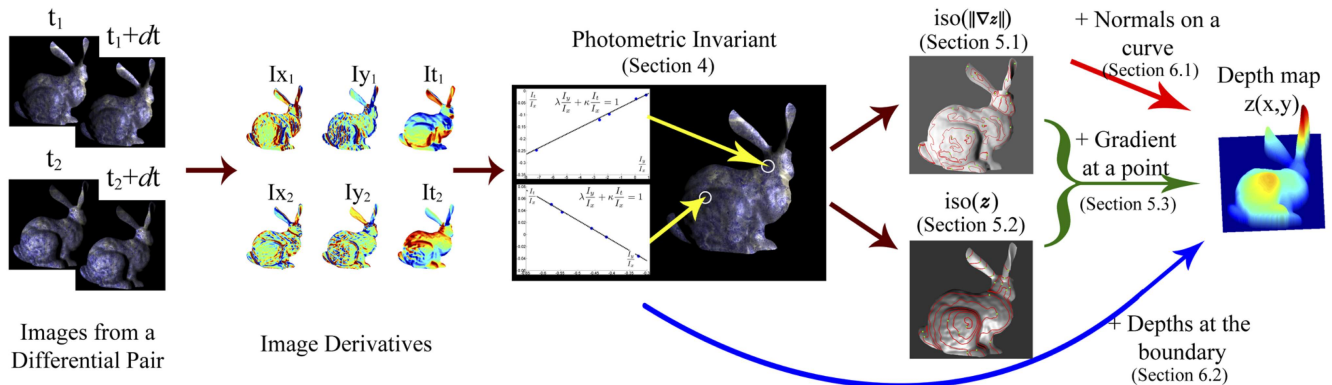


Fig. 1. An outline of our differential theory of photometric stereo. Two or more differential pairs lead to an uncalibrated photometric invariant. Isocontours of constant $\|\nabla z\|$ and constant z may be recovered from the invariant. Additional information in the form of gradient at a single point suffices to recover depths. Normals on a curve or depths on the boundary can be used for a more noise robust estimation.

and isocontours of constant magnitude of the gradient. Further, we show that for general surfaces where these two sets of isocontours intersect transversally, surface normal information at a single point suffices to determine depth over the entire surface. Thus, in theory, differential photometric stereo allows recovery of surface depths from unknown light positions on a circle, for unknown isotropic BRDFs.

While the theory suggests an elegant algorithm for depth reconstruction, recovering isocontours of constant depth requires higher-order derivatives, which may lead to noisy estimation. For practical applications, Section 8 presents two additional algorithms that require only first-order image derivatives to recover surface normals and depths, respectively, given additional information in the form of normals on a curve, or depths at the boundary. Figure 1 outlines our major theoretical results for differential photometric stereo.

Throughout the paper, we validate the theoretical results with several synthetic and real data examples. Implementation details and a simple acquisition setup for differential photometric stereo are presented in Section 9. We highlight the scope of this work with a discussion of future directions in Section 10.

To summarize the above discussion, this paper introduces a comprehensive theory of photometric reconstruction using image derivatives, resting on several fundamental contributions:

- Novel photometric invariants that relate image derivatives to the geometry of a surface with arbitrary isotropic BRDF, without requiring light positions.
- In a shape from shading framework, recovery of isocontours of constant magnitude of gradient.
- In a photometric stereo framework, a photometric flow relates surface geometry directly to image derivatives, with no dependence on light source positions.
- Two differential measurements suffice to extract surface information from the photometric flow relation.
- Recovery of isocontours of constant depth and constant magnitude of the gradient from photometric flow.
- Theoretical recovery of entire surface depth by specifying the surface normal at a single point.
- Practical algorithms for surface normal or depth estimation that require only lower order derivatives.

2 RELATED WORK

This work differs from most prior studies in considering differential information within the context of photometric stereo, with unknown BRDF and light directions. Our theoretical results distinguish this work by establishing a minimal requirement of two differential pairs for 3D reconstruction. This paper extends a preliminary version [4] to present shape from shading and photometric stereo in a combined framework, includes descriptions of implementation and acquisition, as well as greater details of proofs and experiments.

Most prior work in photometric stereo has dealt with Lambertian surfaces, for which the surface normal can be recovered from three images [18]. In some cases, specular removal has been used as a pre-processing step [3].

Hertzmann and Seitz use example-based methods for general BRDFs, but require a reference object of the same material [7]. Helmholtz stereopsis [21] eliminates the BRDF by exploiting reciprocity, but needs active control of both light and camera positions. Tan et al. use color-based separation of the diffuse and specular components of dichromatic reflectance for reconstruction [16]. Further empirical properties of the BRDF can be exploited to recover shape, but require dense coverage of the lighting hemisphere [2], [8].

Isocontours of constant depth are recovered by Alldrin and Kriegman [1] using bilateral symmetry of a spatially varying BRDF. They require a dense configuration of known lights on a circle and initial information on a curve to recover depth. Tan and Zickler [17] further recover isocontours of constant magnitude of the gradient, whereby initial information at a single point suffices. In contrast, our theory does not depend on pointwise BRDF symmetries, rather uses image derivatives. Instead of the dense, calibrated ring of lights of [1], [17], we need as few as two differential pairs, at unknown positions on a circle. While the results of [17] are derived only for the gnomonic projection, our theory holds for many projections (see Sec. 3.1). Thus, we can also recover isocontours of constant depth as well as constant magnitude of the gradient, but our derivations are more general and may lead to other interesting topological classes in 3D reconstruction.

A related work that uses differential information to derive

photometric invariants is Clark's active photometric stereo [5]. However, it recovers depth using a distance-dependent imaging model and calibrated source positions.

Besides photometric stereo, image derivatives have been considered in other contexts. Shape from shading seeks to recover depth from a single image of a Lambertian surface [12], [19]. In computer graphics, Ramamoorthi *et al.* have studied the first-order behavior of reflection as a convolution and visibility under area lighting [14]. Optical flow uses spatial and temporal derivatives for recovering the motion field [11]. Indeed, the form of our photometric invariant bears a striking resemblance to the optical flow constraint. However, our theory neither relies on assumptions like brightness constancy, nor does it suffer from the aperture problem.

3 IMAGE FORMATION FOR ISOTROPIC BRDFS

Throughout this paper, the object and the camera are assumed to be fixed, while a directional point light source may move around the object. The principal axis of the camera is oriented along $\hat{\mathbf{v}} = (0, 0, 1)^\top$, pointing towards the origin. We assume orthographic projection and the object is represented by a surface $z(x, y)$, where $\mathbf{x} = (x, y)^\top$ is a point on the image plane. Unless stated otherwise, a 2-vector on the real plane \mathbb{R}^2 is represented as \mathbf{w} , whereas a unit 3-vector on the 2-sphere $\mathcal{S}^2 \subset \mathbb{R}^3$ is represented as $\hat{\mathbf{w}}$. For a vector $\mathbf{w} \in \mathbb{R}^2$, we denote $l(\mathbf{w}) = \sqrt{\|\mathbf{w}\|^2 + 1}$. The notation f_x stands for the partial derivative of the function f with respect to the variable x .

3.1 A Note on Surface Normal Parameterizations

The relationship between the unit normal $\hat{\mathbf{n}}$ on the 2-sphere and its representation \mathbf{n} is determined by the projection mapping $\pi : \mathcal{S}^2 \rightarrow \mathbb{R}^2$ [9]. For a gnomonic projection, from the center of the sphere to the tangent plane resting on the north pole, $\hat{\mathbf{n}} = \frac{(\mathbf{n}^\top, 1)^\top}{l(\mathbf{n})}$. Also, recall that the unit surface normal, $\hat{\mathbf{n}}$, is related to the surface gradient, $\nabla z = (z_x, z_y)^\top$, as

$$\hat{\mathbf{n}} = \frac{(-z_x, -z_y, 1)^\top}{\sqrt{z_x^2 + z_y^2 + 1}}. \quad (1)$$

Thus, for a gnomonic projection, $\mathbf{n} = [-z_x, -z_y]^\top$. For a stereographic projection centered on the south pole, $\hat{\mathbf{n}} = \frac{(2\mathbf{n}^\top, l(\mathbf{n})^2 - 2)^\top}{l(\mathbf{n})^2}$. Similarly, we represent a directional point light source $\hat{\mathbf{s}}$ by a 2-vector, \mathbf{s} . Figure 2 illustrates the two parameterizations.

For the gnomonic projection, $\mathbf{n} = [-z_x, -z_y]^\top$, so it provides a linear relationship between the surface normal and the gradient. But it maps normals on the occluding contour to infinity, while the stereographic projection is always guaranteed to be finite. For our practical implementations, we choose one of the above two mappings. However, the theory of the paper is independent of this choice and is valid for both projections (and indeed, for any other centered on the line joining the north and south poles).




Fig. 2. An illustration of the two possible parameterizations for the surface normal. The gnomonic projection (red line) is denoted $\mathbf{n} = \pi_g(\hat{\mathbf{n}})$ and the stereographic projection (purple line) is denoted $\mathbf{n} = \pi_s(\hat{\mathbf{n}})$.

3.2 Parameterizing Isotropic BRDFS

We develop our theory for homogeneous isotropic BRDFS. Such reflectance functions depend only on the three angles between the unit normal $\hat{\mathbf{n}}$, the light source direction $\hat{\mathbf{s}}$ and the viewing direction $\hat{\mathbf{v}}$, thus, they can be represented as a function of the form $\bar{\rho}(\hat{\mathbf{n}}^\top \hat{\mathbf{s}}, \hat{\mathbf{n}}^\top \hat{\mathbf{v}}, \hat{\mathbf{s}}^\top \hat{\mathbf{v}})$. For gnomonic projection, the angles that determine the isotropic BRDF are

$$\hat{\mathbf{n}}^\top \hat{\mathbf{s}} = \frac{\mathbf{n}^\top \mathbf{s} + 1}{l(\mathbf{n})l(\mathbf{s})}, \quad \hat{\mathbf{n}}^\top \hat{\mathbf{v}} = \frac{1}{l(\mathbf{n})}, \quad \hat{\mathbf{s}}^\top \hat{\mathbf{v}} = \frac{1}{l(\mathbf{s})}. \quad (2)$$

For a stereographic projection, these angles are given by

$$\begin{aligned} \hat{\mathbf{n}}^\top \hat{\mathbf{s}} &= \frac{4\mathbf{n}^\top \mathbf{s}}{l(\mathbf{n})^2 l(\mathbf{s})^2} + \left(\frac{2}{l(\mathbf{n})^2} - 1 \right) \left(\frac{2}{l(\mathbf{s})^2} - 1 \right), \\ \hat{\mathbf{n}}^\top \hat{\mathbf{v}} &= \frac{2}{l(\mathbf{n})^2} - 1, \quad \hat{\mathbf{s}}^\top \hat{\mathbf{v}} = \frac{2}{l(\mathbf{s})^2} - 1. \end{aligned} \quad (3)$$

Thus, the BRDF can always be written as a function of the form $\hat{\rho}(\mathbf{n}^\top \mathbf{s}, l(\mathbf{n}), l(\mathbf{s}))$, appropriately defined from $\bar{\rho}$:

$$\hat{\rho}(\hat{\mathbf{n}}^\top \hat{\mathbf{s}}, \hat{\mathbf{n}}^\top \hat{\mathbf{v}}, \hat{\mathbf{s}}^\top \hat{\mathbf{v}}) = \hat{\rho}(\mathbf{n}^\top \mathbf{s}, l(\mathbf{n}), l(\mathbf{s})) = \hat{\rho}(\alpha, \beta, \gamma), \quad (4)$$

where, we denote $\alpha = \mathbf{n}^\top \mathbf{s}$, $\beta = l(\mathbf{n})$, $\gamma = l(\mathbf{s})$ for the parameters of an isotropic BRDF. It may be verified that the same is true for a stereographic projection. Note that for the case of photometric reconstruction, the object is fixed with respect to the camera coordinate system, while the distant light source may move across images. Thus, \mathbf{n} depends only on space variables \mathbf{x} , while \mathbf{s} depends only on time variable t .

3.3 Ratio images

The image formation model can be easily generalized to:

$$E(\mathbf{x}, t) = a(\mathbf{x}) \hat{\rho}(\alpha(\mathbf{x}, t), \beta(\mathbf{x}), \gamma(t)), \quad (5)$$

where $a(\mathbf{x})$ is a surface albedo. This is a reasonable model and in particular, subsumes the Lambertian assumption of traditional photometric stereo, while allowing for more general isotropic BRDFS, possibly modulated by a spatially-varying texture or albedo. In fact, we make no assumptions on the form of the isotropic BRDF $\hat{\rho}$.

In practice, one may eliminate the pointwise albedo $a(\mathbf{x})$ by considering ratio images. Note that ratios of arbitrarily

illuminated images need not preserve the functional dependence of the isotropic BRDF on α , β and γ . But ratios with respect to the image obtained from a light source colocated with the sensor, that is, $\mathbf{s} = (0, 0)^\top$, do preserve the desired form of the BRDF. It can be easily seen from (4) that such ratio images, denoted by I , have the form

$$I(\mathbf{x}, t) = \frac{a(\mathbf{x}) \widehat{\rho}(\alpha, \beta, \gamma)}{a(\mathbf{x}) \widehat{\rho}(\beta)} = \rho(\alpha, \beta, \gamma), \quad (6)$$

where $\rho(\cdot)$ is the appropriately defined function. The exact form of $\widehat{\rho}(\cdot)$, $\widehat{\rho}(\cdot)$ or $\rho(\cdot)$ is not important for us, since we will derive our photometric invariant by eliminating it.

Alternatively, one may also take ratios with respect to an image under uniform (floodlit or cloudy sky) illumination.

4 DIFFERENTIAL PHOTOMETRIC INVARIANTS

The key idea that we use to derive our invariants is the linearity of chain rule differentiation. Suppose the image observed at pixel (x, y) at time t is given by $I(x, y, t) = \rho(\mathbf{w})$, where $\mathbf{w} = (w_1, \dots, w_m)^\top \in \mathbb{R}^m$ represents reflectance parameters that depend on geometry, lighting and camera. We assume that the functional forms of the parameters $w_i(x, y, t)$ are known (for instance, lighting direction, half-angle, viewing direction and so on), but the functional form of the BRDF $\rho(\cdot)$ is unknown. Then, the spatial and temporal derivatives are:

$$I_x = \sum_{i=1}^m \rho_{w_i} w_{i,x} \quad (7)$$

$$I_y = \sum_{i=1}^m \rho_{w_i} w_{i,y} \quad (8)$$

$$I_t = \sum_{i=1}^m \rho_{w_i} w_{i,t}. \quad (9)$$

Note that there may be acquisitions or imaging methods where it is possible to have other independent derivatives, not restricted to image plane coordinates (x, y) or a temporal variable (t) .

The linearity of the chain rule leads to a natural factorization of the derivative relations into geometry and BRDF-dependent terms:

$$\begin{bmatrix} I_x \\ I_y \\ I_t \end{bmatrix} = \begin{bmatrix} w_{1,x} & \cdots & w_{m,x} \\ w_{1,y} & \cdots & w_{m,y} \\ w_{1,t} & \cdots & w_{m,t} \end{bmatrix} \begin{bmatrix} \rho_{w_1} \\ \vdots \\ \rho_{w_m} \end{bmatrix}. \quad (10)$$

For convenience, we express this in the form:

$$\nabla \mathbf{I} = \mathbf{G} \nabla \rho \quad (11)$$

Define the matrix $\mathbf{U} = [\mathbf{G} | -\nabla \mathbf{I}]$, then clearly

$$\mathbf{U} [(\nabla \rho)^\top, 1]^\top = \mathbf{0}. \quad (12)$$

Thus, \mathbf{U} is rank-deficient and when it is a square matrix, we can demand the condition

$$\det \mathbf{U} = 0 \quad (13)$$

which is an equation that relates the surface normal, light directions, their derivatives and measured image derivatives.

Importantly, it is a relation which does not depend on the functional form of the BRDF ρ .

The following sections derive such invariants for the particular problems of shape from shading and photometric stereo. While deriving the invariants is straightforward, the real challenge lies in relating them to surface shape in a manner that facilitates reconstruction. As we show throughout the remainder of this paper, careful analysis of the structure of these invariants yields fundamental insights into the extent to which shape may be reconstructed even with a complex, unknown BRDF.

5 SHAPE FROM SHADING

While the primary focus of this paper is on photometric stereo, as a simple initial application, we consider the problem of shape from shading (SFS). Prior work has extensively studied the shape from shading differential equations for the Lambertian case [20]. Extensions to general BRDFs have been considered, but limited to known, parametric reflectance models, such as Torrance-Sparrow [10]. We use our theory to extend the analysis to general, unknown isotropic BRDFs.

It is common to assume in shape from shading frameworks that the light source is colocated with the sensor and the scene albedo is constant [13]. With the source and sensor colocated at $(0, 0, 1)^\top$, note that $l(\mathbf{s}) = 1$ and $\mathbf{n}^\top \mathbf{s} = 0$. Thus, the isotropic BRDF in (4) reduces to a function of the single parameter $\beta = l(\mathbf{n}) = \frac{1}{\sqrt{\|\nabla z\|^2 + 1}}$. Defining $\delta = \|\nabla z\|$, the image formation equation for a colocated source and sensor may equivalently be rewritten as

$$I(x, y) = \rho(\delta(x, y)). \quad (14)$$

The differential image formation equations corresponding to (12) are

$$\begin{bmatrix} \delta_x & -I_x \\ \delta_y & -I_y \end{bmatrix} \begin{bmatrix} \rho_\delta \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad (15)$$

whereby the determinant condition in (13) leads to the relation

$$\frac{I_x}{I_y} = \frac{\delta_x}{\delta_y} = \frac{\|\nabla z\|_x}{\|\nabla z\|_y}. \quad (16)$$

While (16) represents a simple first-order linear PDE in δ , it will be useful for subsequent results in this paper to characterize the solution by the following proposition:

Proposition 1. *From a single image of a surface, obtained using a colocated source and sensor, one may recover isocontours of constant magnitude of the surface gradient.*

Proof: For a surface $z(x, y)$, consider the scalar field $g(x, y) = \|\nabla z(x, y)\|$. The gradient associated with this field is given by $\nabla g = (\|\nabla z\|_x, \|\nabla z\|_y)^\top$. It is well-known that the level curves $g(x, y) = c$, for constant c , are orthogonal to the gradient ∇g . The direction of the tangent to the level curves is $\frac{\|\nabla z\|_y}{\|\nabla z\|_x}$, which is the same as $\frac{I_y}{I_x}$, from (16). Thus, from a single image, one may trace the isocontours of constant $\|\nabla z(x, y)\|$, using the knowledge of their tangent directions. \square




Fig. 3. Illustration of Proposition 1 on synthetic data corresponding to (a) a vase and (b) a Mozart figurine. Given a single image from a sensor collocated with a light source, isocontours of constant $\|\nabla z\|$ are computed. All points on the red curves have the same depth as the corresponding blue dot. The ground truth depth map is shown for reference in the inset.

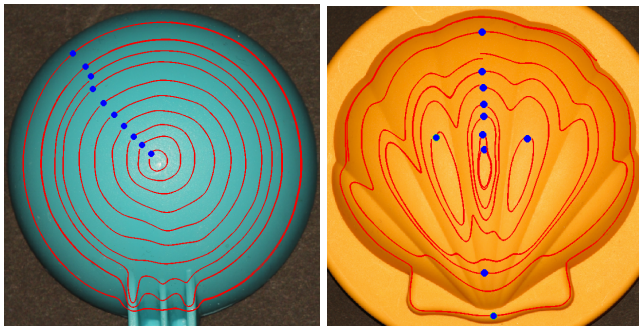


Fig. 4. Given a single real image of (a) a plastic spoon and (b) a plastic shell, the result of Proposition 1 is used to plot isocontours of constant $\|\nabla z\|$, shown in red, starting from the points marked in blue. No radiometric calibration of the input image is required.

5.1 Experimental Evaluation

We illustrate the result of Proposition 1 on synthetic and real data. In Figure 3, we show the recovered isocontours of constant $\|\nabla z\|$ on synthetic vase and Mozart data obtained from [20].

For real data experiments, we place a light bulb close to the camera to approximate a collocated source and sensor. Note that no radiometric calibration is required, since the camera response function is subsumed in the BRDF function ρ . In Figure 4, the recovered isocontours of constant $\|\nabla z\|$ are shown on real images of a plastic spoon and a plastic shell.

6 PHOTOMETRIC STEREO

The remainder of this paper deals with photometric reconstruction using images of a stationary object, acquired with a fixed orthographic camera, under multiple illumination conditions. This problem is commonly referred to as photometric stereo. Allowing motion of the light source (as opposed to shape from shading, where we have a single, fixed light source), makes the images dependent on an additional temporal variable (such as time, or an arc length parameter if motion is along a curve).

A common approach to recovering shape from photometric information is to derive invariants that relate the image intensities to surface geometry and light source directions. Most prior work focuses on more restricted classes of BRDF, such as a Lambertian one. However, one may deal with complex

BRDFs if they can be eliminated from a system of equations, as outlined in our general framework of Section 4. In this section, we use differential information to derive such an invariant. A surprising and useful result is that *our invariant does not depend on light source positions*, so it is in fact an uncalibrated invariant.

6.1 Differential Images

Intuitively, a single image formation equation leads to independent relations upon differentiation with respect to various space or time variables. These may then be related by eliminating any terms that depend on the functional form of the BRDF. In this section, we use this intuition to derive a novel photometric invariant.

The space and time derivatives of the images in (6) are

$$\nabla_x I(\mathbf{x}, t) = \rho_\alpha \mathbf{J}^\top(\mathbf{n}) \mathbf{s} + \rho_\beta \frac{1}{l(\mathbf{n})} \mathbf{J}^\top(\mathbf{n}) \mathbf{n} \quad (17)$$

$$I_t(\mathbf{x}, t) = \rho_\alpha \mathbf{s}_t^\top \mathbf{n} + \rho_\gamma \frac{1}{l(\mathbf{s})} \mathbf{s}_t^\top \mathbf{s} \quad (18)$$

where α, β, γ are defined previously and $\mathbf{J}(\mathbf{n}) = [\mathbf{n}_x, \mathbf{n}_y]$ is the 2×2 Jacobian related to the second fundamental form, II . (For the gnomonic projection, $\mathbf{J}(\mathbf{n}) = l(\mathbf{n}) \cdot II$.)

This system of three equations (note that (17) represents two equations) is clearly underconstrained, with the unknowns $\rho_\alpha, \rho_\beta, \rho_\gamma, \mathbf{n}$ and \mathbf{s} . However, one way to extract constraints on the normal might be to eliminate the BRDF derivatives, which can be done from these three equations if one of ρ_α, ρ_β or ρ_γ terms can be made to identically vanish. This can be achieved in a setup of circular motion, as discussed next.

6.2 Circular Motion Yields an Invariant

Let us constrain the source to move in a circle around the camera axis. Intuitively, since $\|\mathbf{s}\|$ is now constant, γ stays constant. Thus, the isotropic BRDF reduces to a 2D function, which allows elimination of derivatives with respect to α and β . Mathematically, $\mathbf{s}_t^\top \mathbf{s} = 0$ for lights on a circle, so the system of three differential equations in (17) and (18) reduces to a form analogous to (12):

$$\begin{bmatrix} \mathbf{s}^\top \mathbf{n}_x & \frac{1}{l(\mathbf{n})} \mathbf{n}^\top \mathbf{n}_x & -I_x \\ \mathbf{s}^\top \mathbf{n}_y & \frac{1}{l(\mathbf{n})} \mathbf{n}^\top \mathbf{n}_y & -I_y \\ \mathbf{s}^\top \mathbf{n} & 0 & -I_t \end{bmatrix} \begin{bmatrix} \rho_\alpha \\ \rho_\beta \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \quad (19)$$

The determinant condition of (13) leads to the relation

$$\frac{I_x \mathbf{s}_t^\top \mathbf{n} - I_t \mathbf{s}^\top \mathbf{n}_x}{I_y \mathbf{s}_t^\top \mathbf{n} - I_t \mathbf{s}^\top \mathbf{n}_y} = \frac{\mathbf{n}^\top \mathbf{n}_x}{\mathbf{n}^\top \mathbf{n}_y}. \quad (20)$$

Clearly, the right hand side depends only on position, while the left hand side contains time-dependent entities. Thus, a photometric invariant that relates image derivatives, surface geometry and light directions can be expressed as

$$\frac{I_x \mathbf{s}_t^\top \mathbf{n} - I_t \mathbf{s}^\top \mathbf{n}_x}{I_y \mathbf{s}_t^\top \mathbf{n} - I_t \mathbf{s}^\top \mathbf{n}_y} = \text{constant across time.} \quad (21)$$

6.3 The Invariant Is Independent of Light Positions

It might seem at a first glance that, given light source positions, a non-linear minimization framework can be used to solve for the unknowns $\{\mathbf{n}, \mathbf{n}_x, \mathbf{n}_y\}$, up to scale, using the constraint in (20). However, as we show below, the light source directions are not required at all. That is, the invariant (21) directly relates image derivatives to surface geometry.

Proposition 2. *The entire information in the photometric invariant (21) is encapsulated by two entities, which depend only on surface geometry and not on source positions \mathbf{s} .*

While the above may seem surprising given the presence of light source terms in (21), a closer look at the structure of (21) immediately proves the above proposition.

Proof: Defining

$$\lambda = \frac{\mathbf{n}^\top \mathbf{n}_x}{\mathbf{n}^\top \mathbf{n}_y}, \quad \mathbf{u} = \mathbf{n}_x - \lambda \mathbf{n}_y, \quad (22)$$

we can rewrite (20) as

$$\frac{I_x - \lambda I_y}{I_t} = \frac{\mathbf{s}^\top \mathbf{u}}{\mathbf{s}_t^\top \mathbf{n}} = \kappa. \quad (23)$$

By definition, λ depends only on the surface normal and is constant across time (that is, independent of light source positions). Also, by definition of λ and \mathbf{u} in (22), we have $\mathbf{u}^\top \mathbf{n} = 0$, that is, \mathbf{u} is orthogonal to \mathbf{n} . Since the lights lie on a circle, \mathbf{s}_t is also orthogonal to \mathbf{s} . Moreover, at time t (or equivalently, angular position on the circle of sources), the light source is $\mathbf{s} = (r \cos t, r \sin t)^\top$. Then, $\mathbf{s}_t = (-r \sin t, r \cos t)^\top$, thus, $\|\mathbf{s}\| = \|\mathbf{s}_t\|$. It follows that κ is actually the (signed) ratio of the magnitudes of \mathbf{u} and \mathbf{n} .¹

Indeed, with $\mathbf{n}_\perp = (-n_2, n_1)^\top$, it immediately follows:

$$\mathbf{u} = -\kappa \mathbf{n}_\perp. \quad (24)$$

Thus, we have the following constraint, equivalent to the invariant in (20), but which *does not depend on knowledge of light source positions*:

$$I_x - \lambda I_y - \kappa I_t = 0. \quad (25)$$

Clearly, all the information in the invariant is encapsulated by λ and κ , which depend only on surface geometry. \square

6.4 Importance of the Invariant

The result of Proposition 2 is a surprising one and in fact, can be understood as a fundamental relationship between spatial and temporal derivatives of images arising from isotropic BRDFs:

Corollary 3. *For a surface with isotropic BRDF, the $N \times 3$ matrix of spatial and temporal image derivatives at a pixel, recorded for $N > 1$ unknown light positions on a circle, must be rank 2. In addition, the null-vector, denoted $(\lambda, \kappa, 1)^\top$, depends only on the surface geometry.*

1. In practice, the temporal derivative is obtained as a difference between images at time t and $t + \delta t$, so asserting $\|\mathbf{s}\| = \|\mathbf{s}_t\|$ assumes that the angular difference δt between the lights of the rotating differential pair is known. This is the same as the ratio of their distance and the radius of the circle on which they are situated. Note that just an unmarked piece of string may suffice to create any integral ratio.

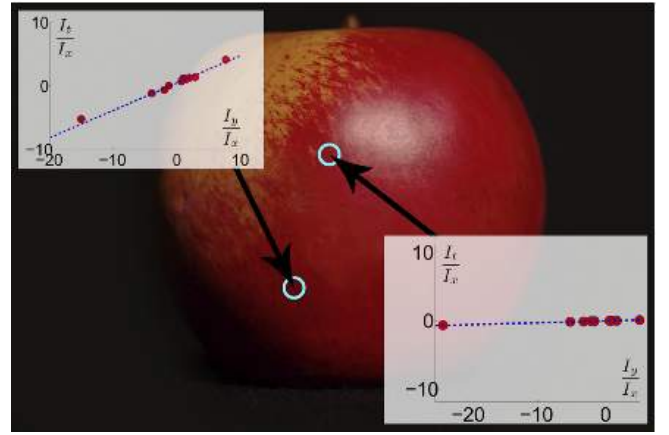


Fig. 5. The ratio of image derivatives at a pixel, recorded for various light positions, lie on a straight line given by the equation $\lambda \frac{I_y}{I_x} + \kappa \frac{I_t}{I_x} = 1$.

The above is also useful practically, since it raises the possibility that surface information may be recovered from image derivatives in photometric stereo for isotropic BRDFs, without knowledge of light source positions. In fact, as few as two differential image pairs suffice to estimate λ and κ by solving a small linear system. This is an important observation:

Corollary 4. *Two pairs of differential images suffice to recover surface information contained in the invariant of (21).*

Note that three light sources can create two differential pairs. It is instructive to recall that traditional photometric stereo for Lambertian surfaces requires three images to completely determine the surface normal.

At this stage, it is natural to seek a characterization of the exact surface information recoverable from λ and κ . In the following section, we will show that knowledge of λ and κ , together with initial information at a single point, suffices to completely determine the surface depths.

6.5 Experimental Evaluation

Here, we empirically illustrate the validity of the relation in (25). A differential pair of lights is moved on a circle and real images of a plastic apple of varying albedo are acquired at 11 unknown light positions. Ratio images are computed with respect to a floodlit image.

Figure 5 illustrates that the ratios $\frac{I_y}{I_x}$ and $\frac{I_t}{I_x}$ computed for the 11 different light positions lie close to a straight line. The entities λ and κ are given by the best-fit straight line.

7 PHOTOMETRIC FLOW

It is tempting to compare the form of equation (25) to the optical flow relation [11]. Thus, we refer to (25) as the *photometric flow* relation. However, we note that the photometric flow relation is derived without resorting to a brightness constancy assumption and as established by Corollary 4, recovery of λ and κ in a photometric stereo setup does not suffer from the aperture problem.

The previous section demonstrates that the entire information in the photometric flow can be expressed (and recovered) in terms of the two entities λ and κ . However, it is not immediately clear what that means in terms of surface reconstruction. We answer that conclusively in the present section.

7.1 Isocontours of Constant $\|\nabla z\|$

To begin, we show that one may recover isocontours where magnitude of the gradient, $\|\nabla z\|$, stays constant.

Corollary 5. *From two or more differential images of a surface, obtained from unknown light source positions, one may recover isocontours of constant magnitude of the surface gradient.*

Proof: This follows directly from the definition of λ in Proposition 2:

$$\lambda = \frac{\mathbf{n}^\top \mathbf{n}_x}{\mathbf{n}^\top \mathbf{n}_y} = \frac{\|\mathbf{n}\|_y}{\|\mathbf{n}\|_x}. \quad (26)$$

The remainder of the proof follows that of Proposition 1. \square

In the context of photometric stereo, prior work by Tan and Zickler has recovered isocontours of constant $\|\nabla z\|$, but relying on properties of the gnomonic projection [17]. In contrast, note that our proof is actually more general than the statement of Proposition 5. It may be easily verified by the reader that the isocontours of constant $\|\mathbf{n}\|$ are the same for any \mathbf{n} derived from a projection centered on the line joining the poles. In the context of the gnomonic projection where $\mathbf{n} = \nabla z$, isocontours of constant $\|\mathbf{n}\|$ are called *isoslope contours* in [17].

7.1.1 Experimental Evaluation

In Figure 6, we illustrate the recovery of isocontours of constant $\|\nabla z\|$ using synthetic data. For the images of a synthetic sphere of varying albedo and the synthetic bunny, we use a simplified Torrance-Sparrow model

$$\bar{\rho} = \frac{1}{4\pi\sigma^2} \exp \left[- \left(\frac{\cos^{-1} \hat{\mathbf{n}}^\top \hat{\mathbf{h}}}{\sigma} \right)^2 \right], \quad \hat{\mathbf{h}} = \frac{\hat{\mathbf{s}} + \hat{\mathbf{v}}}{\|\hat{\mathbf{s}} + \hat{\mathbf{v}}\|}. \quad (27)$$

with $\sigma = 0.3$. For images of the vase of varying albedo, we use a constant coefficient Phong-Blinn model: $\bar{\rho} = \hat{\mathbf{n}}^\top \hat{\mathbf{s}} + (\hat{\mathbf{n}}^\top \hat{\mathbf{h}})^\sigma$, with a typical value of $\sigma = 5$.

In each case, we observe that the recovered isocontours of constant $\|\nabla z\|$ match the ground truth very closely.

7.2 Isocontours of Constant Depth

To recover isocontours of constant $\|\mathbf{n}\|$, we required knowledge only of λ . However, we should be able to further disambiguate the surface using κ . The following proposition gives a constructive proof that λ and κ together determine the isocontours of constant depth (in addition to the isocontours of constant $\|\nabla z\|$).

Proposition 6. *From two or more differential images of a surface, obtained from unknown source positions, one may recover the direction of the surface gradient at every point.*

Proof: Let us denote $\mathbf{n} = (p, q)^\top$, thus, $\mathbf{n}_x = (p_x, q_x)^\top$ and $\mathbf{n}_y = (p_y, q_y)^\top$. Given two or more pairs of differential images, one may estimate $\lambda(x, y)$ and $\kappa(x, y)$ using the linear relation in (25). Thereby, one obtains two linear, first order PDEs from (24) at every pixel:

$$p_x - \lambda(x, y)p_y = \kappa(x, y)q \quad (28)$$

$$q_x - \lambda(x, y)q_y = -\kappa(x, y)p. \quad (29)$$

Further, since we are dealing with a surface, it must satisfy the integrability condition:

$$p_y - q_x = 0. \quad (30)$$

Thus, we have a coupled first order system of three linear PDEs in the two variables p and q . Note that this is an overdetermined system, which may not be solvable in general. However, in our particular case, we can exploit the special form of the PDEs to derive an unusual solution.

Consider the function $h = p - \lambda q$. Taking partial derivatives, we get

$$\begin{aligned} h_x &= p_x - \lambda q_x - \lambda_x q \\ &= p_x - \lambda p_y - \lambda_x q = (\kappa - \lambda_x)q \end{aligned} \quad (31)$$

$$\begin{aligned} h_y &= p_y - \lambda q_y - \lambda_y q \\ &= q_x - \lambda q_y - \lambda_y q = -\kappa p - \lambda_y q. \end{aligned} \quad (32)$$

By integrability of h , we have $h_{xy} = h_{yx}$, which gives us another first order linear PDE:

$$(\lambda_x - \kappa)q_y - \lambda_y q_x - \kappa p_x = \kappa_x p + \kappa_y q. \quad (33)$$

Thus, from (28), (29), (30) and (33), we have a linear system in $\{p_x, p_y, q_x, q_y\}$, given by

$$\begin{bmatrix} 1 & -\lambda & 0 & 0 \\ 0 & 0 & 1 & -\lambda \\ 0 & 1 & -1 & 0 \\ -\kappa & 0 & -\lambda_y & \lambda_x - \kappa \end{bmatrix} \begin{bmatrix} p_x \\ p_y \\ q_x \\ q_y \end{bmatrix} = \begin{bmatrix} \kappa q \\ -\kappa p \\ 0 \\ \kappa_x p + \kappa_y q \end{bmatrix}, \quad (34)$$

whereby expressions for $\{p_x, p_y, q_x, q_y\}$ are obtained in terms of linear functions of $\{p, q\}$:

$$p_x = \nu_1 p + \eta_1 q, \quad p_y = \nu_2 p + \eta_2 q \quad (35)$$

$$q_x = \nu_2 p + \eta_2 q, \quad q_y = \nu_3 p + \eta_3 q. \quad (36)$$

with

$$\begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix} = \frac{1}{\Delta} \begin{bmatrix} \lambda(\kappa^2 - \lambda_x \kappa + \lambda \kappa_x) \\ \kappa^2 - \lambda_x \kappa + \lambda \kappa_x \\ -(\lambda \kappa^2 + \lambda_y \kappa - \kappa_x) \end{bmatrix}, \quad (37)$$

$$(38)$$

$$\begin{bmatrix} \eta_1 \\ \eta_2 \\ \eta_3 \end{bmatrix} = \begin{bmatrix} \kappa \\ 0 \\ 0 \end{bmatrix} + \frac{\kappa_y + \kappa^2}{\Delta} \begin{bmatrix} \lambda^2 \\ \lambda \\ 1 \end{bmatrix}, \quad (39)$$

where $\Delta = -(\kappa \lambda^2 + \lambda \lambda_y - \lambda_x + \kappa)$ is the (3, 2)-minor of the 4×4 matrix in (34). At this stage, the reader may verify by substitution that the expressions above are consistent with the definitions of λ and κ in (22).




Fig. 6. Recovery of isocontours of constant $\|\nabla z\|$. Red curves plot the isocontours, while the green dots represent points chosen to start tracing the curves. (a) A sphere with variable albedo and simplified Torrance-Sparrow BRDF. (b) A vase with Blinn-Phong BRDF. (c) A bunny with Torrance-Sparrow BRDF. (d) Ground truth isocontours for the bunny.

Now, we again use integrability of p and q , that is $(p_x)_y = (p_y)_x$ and $(q_x)_y = (q_y)_x$, to get two new PDEs:

$$(\nu_1 p + \eta_1 q)_y = (\nu_2 p + \eta_2 q)_x \quad (40)$$

$$(\nu_2 p + \eta_2 q)_y = (\nu_3 p + \eta_3 q)_x, \quad (41)$$

These are linear, first order PDEs, where we can again replace the first order derivatives $\{p_x, p_y, q_x, q_y\}$ using (35) and (36) to get two (dependent) homogeneous linear equations in $\{p, q\}$. Using, say, the first equation, we get the ratio of p and q as

$$\frac{q}{p} = \frac{\nu_{1y} + \eta_1 \nu_3 - \nu_{2x} - \nu_2 \eta_2}{\nu_2 \eta_1 + \eta_2 x + \eta_2^2 - \nu_1 \eta_2 - \eta_1 y - \eta_1 \eta_3} \quad (42)$$

The statement of the theorem follows by considering a gnomonic projection, where $p = -z_x$ and $q = -z_y$. \square

Again, we note that the actual result proved is stronger and holds for several projections besides gnomonic. To emphasize the import of the result: using just Gaussian elimination and repeated use of the integrability condition, just two pairs of differential images at unknown light source positions allow us to recover the direction of the gradient at every point of a surface with isotropic BRDF, without requiring any additional information.

Similar to Corollary 5, it immediately follows that:

Corollary 7. *From two or more differential images of a surface, obtained from unknown source positions, it is possible to recover the isocontours of constant depth (or the level curves) for the entire surface.*

This result may be contrasted with symmetry-based methods [1] that theoretically require a dense configuration of lights at known positions to recover the same information.

7.2.1 Experimental evaluation

The recovery of the direction of the gradient and the level curves of constant depth is illustrated for synthetic data in Figure 7. To display the direction of the gradient, the arrows are plotted with length normalized to one. The isocontours of constant depth are shown in red. All points on a red curve have the same depth as the green point on that curve.




Fig. 7. Recovery of isocontours of constant depth. (Top row) Red curves plot the isocontours with the same depth as the green dots. (Bottom row) Plots of the direction of the gradient, with the length of the vector normalized to one for display. (a) A vase with variable albedo and Blinn-Phong BRDF. (b) A bunny with Torrance-Sparrow BRDF.

7.3 Surface Reconstruction from Isocontours

We note that one may not recover magnitude of the gradient without additional information, since the uncalibrated invariant of (21) is homogeneous in p , q and their derivatives. The following proposition establishes that, with the results of Corollaries 5 and 7, additional information is required for general surfaces only on a set of measure zero.

We assume that the surface can be split into a finite number of regions, each of which satisfies a generality condition, namely, that the isocontours of constant z and constant $\|\nabla z\|$ are not everywhere parallel. Surfaces such as a hemisphere are not general, but most surfaces do satisfy this condition.

Proposition 8. *Given the isocontours of constant depth and constant $\|\nabla z\|$ for a general surface $z(x, y)$, specification of the surface normal at a single point suffices to recover depth up to a global convex-concave ambiguity and additive offset.*




Fig. 8. Transversality of isocontours of constant z and constant $\|\nabla z\|$ is sufficient to recover depth, given the surface normal at a single point \mathbf{x}^* . See Proposition 8.

Proof: Let \mathbf{x}^* be a point in an open set \mathcal{U} where the isocontours of constant z and constant $\|\nabla z\|$ intersect transversally. Then, within \mathcal{U} , one may define unit vector fields \mathbf{v} and \mathbf{w} that are tangent, respectively, to the isocontours of constant z and constant $\|\nabla z\|$ and thus, are transversal. Then, since ∇z is orthogonal to \mathbf{v} , it is apparent that the following relation must hold at the point $\mathbf{x}^* \in \mathcal{U}$ (see Figure 8)

$$\mathbf{w} \cdot \frac{\nabla z}{\|\nabla z\|} = \pm \sqrt{1 - (\mathbf{v} \cdot \mathbf{w})^2}. \quad (43)$$

Note the sign ambiguity, which arises since the directions of \mathbf{v} and \mathbf{w} can be specified at most up to a global sign. If the value of $\|\nabla z\|$ is specified at \mathbf{x}^* , it is also specified at every point on the isocontour of constant $\|\nabla z\|$. Thus, we have a linear ODE in z along that isocontour, which may be solved up to an additive constant, c^* .

Since the isocontours of constant z and constant $\|\nabla z\|$ are transversal in \mathcal{U} , the values of depths can now be assigned, up to an unknown c^* , along all the isocontours of constant z that intersect the isocontour of constant $\|\nabla z\|$ passing through \mathbf{x}^* and thus, almost everywhere on \mathcal{U} .

Note that the sign ambiguity in (43) corresponds to a global convex-concave ambiguity and the unknown constant c^* corresponds to a global additive offset.

Finally, we note that for a unit normal $\hat{\mathbf{n}} = (n_1, n_2, n_3)^\top$, we have $\|\nabla z\| = \sqrt{\left(\frac{1}{n_3}\right)^2 - 1}$. Thus, specifying the surface normal at a single point is sufficient to determine $\|\nabla z\|$ and seed the above depth reconstruction. \square

Again, we contrast with calibrated methods [1], which use dense sources to recover isocontours of constant z and require additional information on an entire curve to resolve the depth. As the above proof shows, incorporating gradients has the advantage of reducing this ambiguity, in an uncalibrated framework with only two differential images, to the specification of information at a single point.




Fig. 9. Depth recovery for a monkey saddle surface. (a) The ground truth surface. (b) Recovered isocontours of constant $\|\nabla z\|$ (red) and constant z (cyan), from images under a Torrance-Sparrow BRDF model. All points on the red and cyan curves have same depths as the respective green and magenta points. (c) Specifying the surface normal at one point (marked in blue) on the surface determines all depths along the corresponding isocontour of constant $\|\nabla z\|$ (green points), from which depth at every other point can be determined by tracing isocontours of constant z (red curves). (d) The reconstructed surface by tracing all the isocontours.

7.3.1 Experimental Evaluation

To empirically demonstrate reconstruction from isocontours, we simulate a monkey saddle, $z = x^3 - 3xy^2$, for which the isocontours of constant z and $\|\nabla z\|$ are transversal. In Fig. 9(b), we show recovery of isocontours of constant z and $\|\nabla z\|$. Specifying the normal at a single point on one of the isocontours of constant $\|\nabla z\|$ allows us to assign depths to isocontours of constant z (Fig. 9(c)) and recover the depth map (Fig. 9(d)). The recovered depth is nearly the same as ground truth, except in regions where the chosen isocontour (dotted green in Fig. 9(c)) does not cross level curves of z .

8 PRACTICAL ALGORITHMS

At this stage, we recall that given image derivatives, one may estimate the entities λ and κ that contain sufficient information for disambiguating the surface, given the surface normal at a single point. However, in practice, the equation (42) involves third-order derivatives of the surface normal, which can lead to noisy estimation.

One alternative to recover the surface normals or depth is to incorporate additional information in the form of initial or boundary conditions. In this section, we present algorithms for recovering surface normals given initial information across a curve, or recovering depths given boundary information. These algorithms require only the estimation of λ and κ , which can be performed directly from image derivatives, without resorting to higher order differentiation.

This section focuses on presenting the reconstruction algorithms and their results on synthetic and real data. The experimental details of acquisition and implementation are presented in Section 9.




Fig. 10. Reconstruction for a synthetic sphere using Algorithm 1 of Section 8.1. (a) Isocontours of constant $r = \|\nabla z\|$ computed using the estimated $\lambda(x, y)$, followed by estimation of $\theta = \tan^{-1}\left(\frac{z_x}{z_y}\right)$ using seed values of the surface normal specified on the brown dots. The values of θ are shown using a color-code that maps the interval $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$ between blue and red. (b) Recovered values of θ , with the specified color map. (c) Ground truth θ . (d) Estimated gradient in the x direction, z_x , red indicates large positive value and blue indicates large negative value. (e) Estimated gradient in the y direction, z_y , with the same color map. (f) Side view of the height recovered from integration of estimated z_x and z_y , which matches the ground truth shape. Red indicates higher heights and blue indicates lower heights.

8.1 Algorithm 1: Recovering Surface Normals

Proposition 9. *Initial data in the form of known surface normals on a curve suffices to recover the surface normals from the coupled PDEs in (28) and (29).*

Note that including the integrability requirement of (30) is ignored here (as is done in traditional Lambertian photometric stereo). Once the surface normals are estimated, one may impose integrability while recovering the depth map.

Proof: Consider the isocontours of constant $\|\nabla z\|$, parameterized by x , which are solutions to the ODE $\dot{y} = -\lambda(x, y)$. For a closed, smooth surface, these characteristic curves will, in general, be non-intersecting.

Along the above curves, the pair of PDEs in (28) and (29) reduces to a pair of ODEs:

$$\dot{p} = \kappa q, \quad \dot{q} = -\kappa p. \quad (44)$$

Let $p = r \sin \theta$ and $q = r \cos \theta$. Then, the following pair of relations, obtained by differentiating p and q with respect to the curve parameter (in this case, x) and substituting in the above pair of ODEs, must be true:

$$\dot{r} \cos \theta = (r \sin \theta)(\dot{\theta} - \kappa), \quad \dot{r} \sin \theta = (r \cos \theta)(\kappa - \dot{\theta}). \quad (45)$$

Thus, $\dot{r}^2 = r^2(\dot{\theta} - \kappa)^2$. So, a solution to the pair of ODEs can be obtained as

$$\dot{r} = 0, \quad \dot{\theta} = \kappa. \quad (46)$$

Given initial data along a curve, we can solve the above pair of ODEs. These initial conditions amount to specifying r and θ along a curve. \square

8.1.1 Experimental Evaluation

We demonstrate the applicability of Algorithm 1 using synthetic and real data. For a synthetic sphere with arbitrary isotropic BRDF, five differential pairs of images at unknown light source positions are used to estimate the photometric flow. Given the estimated values of $\lambda(x, y)$, Figure 10(a) shows a few contours of constant r (that is, the curves $\dot{r} = 0$), traced with the initial seed point marked by a brown dot. At the brown dots, we specify the value of the surface normal, thus, the value of θ .

Then, along the isocontours, we solve the integral equation $\theta = \int \kappa dx$, with the constant of integration being the specified value at the brown dot. The recovered values of θ are shown color-coded on the corresponding isocontours in Figure 10(a).

Figure 10(b) shows the values of θ recovered by tracing isocontours over all the pixels of the image. The artifacts near the center are due to discretization (our isocontours are traced with sub-pixel precision, but the values of θ can be displayed only up to a 1 pixel precision). For comparison, we show the ground truth value of θ in Figure 10(c).

Given the values of r and θ , we can recover the gradients z_x and z_y using the above equations. The recovered values are shown in Figures 10(d) and (e). Finally, using the recovered gradients, one may perform a surface normal integration to estimate the depth map, shown in Figure 10(f).

Note that the algorithm requires knowledge of surface normals along the curve given by the brown dots in Figure 10(a), which is difficult to obtain in practice. But if initial data on a curve is indeed available, Algorithm 1 is an elegant method to recover the surface normals directly from first order image derivatives, without resorting to higher order differentiation.

In Figure 11, we show the recovered isocontours of constant $\|\nabla z\|$ on real data for a hand-moulded clay ball and a plastic apple. If the surface normals are specified at a single point on these contours, the shape may be recovered. Note that one may not use the occluding contour as the initial curve, as it can be shown to be non-transversal to isocontours of constant $\|\mathbf{n}\|$. (At the occluding contour, since $\hat{\mathbf{n}}^\top \hat{\mathbf{v}} = 0$, it follows from (1) that $\|\mathbf{n}\| \rightarrow \infty$ for viewing direction $\hat{\mathbf{v}} = (0, 0, 1)^\top$.)

8.2 Algorithm 2: Recovering Depth

Proposition 10. *Boundary data in the form of known depths suffices to recover the surface depth from the PDEs in (28), (29) and (30).*

Proof: Using the integrability condition (30), the pair of equations (28) and (29) can be written as a single constraint:

$$p_x = \lambda^2 q_y - \lambda \kappa p + \kappa q. \quad (47)$$

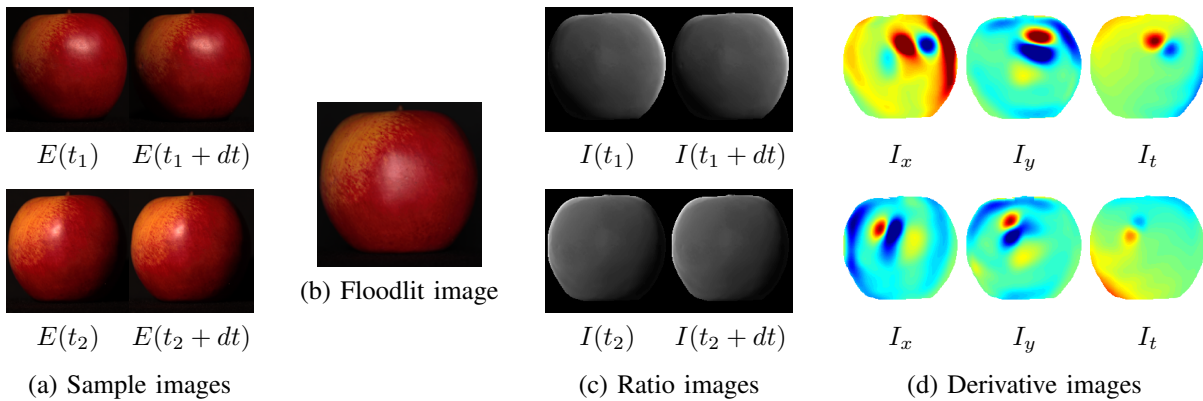


Fig. 12. Input pipeline for the apple dataset. (a) Sample images from the differential pair. (b) Floodlit image. (c) Ratios with respect to the floodlit image. (d) Spatial and temporal derivatives. Blue indicates low values and red indicates high values.




Fig. 11. Isocontours of constant $\|\nabla z\|$ for real data corresponding to a hand-moulded clay ball and a plastic apple.

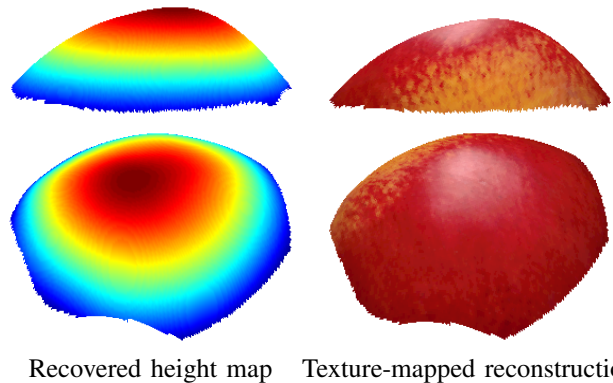


Fig. 13. Reconstruction on real data for the apple dataset using Algorithm 2 of Section 8.2, with Dirichlet boundary conditions. (Top row) Side view. (Bottom row) Top and side view. (Left column) Height map, red indicates higher values and blue stands for lower values. (Right column) Texture-mapped reconstruction.

For a gnomonic projection, we have $p = -z_x$ and $q = -z_y$. Thus, we can rewrite the above as

$$z_{xx} - \lambda^2 z_{yy} + \lambda \kappa z_x - \kappa z_y = 0. \tag{48}$$

This is a linear, second-order hyperbolic PDE, which is well-posed given Dirichlet boundary conditions. \square

While demanding the satisfiability of the constraint in (48) is, in fact, a weaker condition than the coupled constraints in (28) and (29), it leads to a convenient numerical implementation. In practice, depths may be specified at the boundary for scenes where an object rests on a background plane. Indeed, the extensive theory of solutions for hyperbolic PDEs can be used to solve the PDE in (48).

8.2.1 Experimental Evaluation

In Figure 12(a), we show 2 of the 11 differential pair of images for a plastic apple. Note that the object has variable albedo and a non-Lambertian BRDF. These input images are divided by the floodlit image in Figure 12(b) to obtain the ratio images in Figure 12(c). Note that the albedo variations are eliminated in the ratio images. Spatial and temporal derivatives are computed on the ratio images (Figure 12(d)). From these derivatives, we perform a reconstruction using Algorithm 2 of this section, with boundary depths set to 0. As seen in Figure 13, the 3D reconstruction closely resembles the expected shape.

In Figure 14, as a means of comparison to ground truth, we show experimental results for a real teflon sphere with uniform albedo. Note that traditional Lambertian photometric stereo can

be performed with just two light sources, or several coplanar lights, for an object with uniform albedo. Figure 14(a) shows 2 of the 13 differential pairs used for reconstruction, while Figure 14(b) shows the corresponding spatial and temporal derivatives. It can be easily seen from the images that the material of the sphere is non-Lambertian. Consequently, the height map reconstructed by traditional photometric stereo using 13 lights is clearly sheared (Figure 14(c)). In contrast, the reconstructed height map using the algorithm presented in this section closely resembles a sphere (Figure 14(d)).

In Figure 15(a), we show input images for a toy dog, painted with water color. Note the coarse and fine scale variations in the albedo, as well as some non-Lambertian effects. More importantly, the object surface is not differentiable, so it presents a challenging scenario for the theory of this paper. We again acquire a floodlit image (Figure 15(b)) and compute ratio images to eliminate the albedo and obtain images whose intensities depend only on the BRDF (Figure 15(c)).

The reconstruction using Algorithm 2, with boundary depths set to 0, is shown in Figure 16(a). A depth map for the reconstruction is shown in Figure 16(b). A texture-mapping on the same reconstruction is shown in Figure 16(c). In Figure 16(d), we show a close-up of the reconstruction for the head.

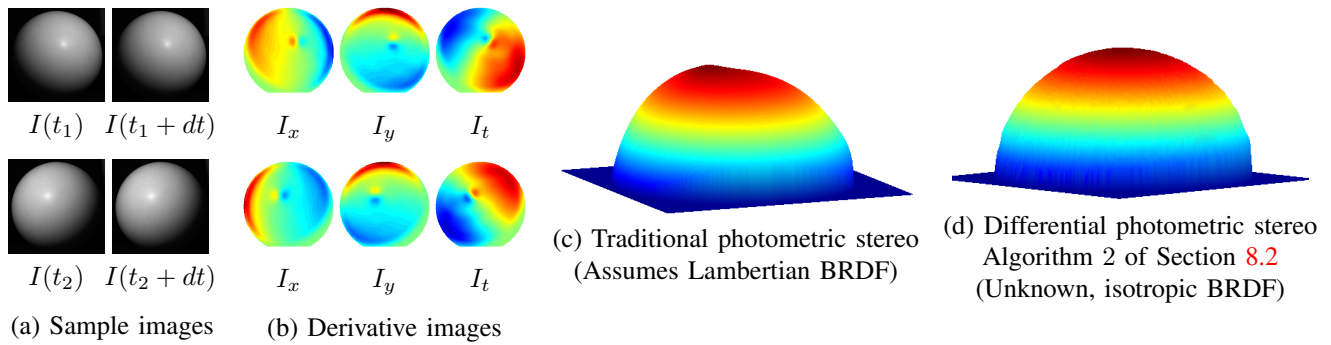


Fig. 14. Comparison with traditional Lambertian photometric stereo, using real data. (a) Two sample differential pairs of images of a teflon ball. (b) Spatial and temporal image derivatives. Red stands for positive values and blue stands for negative values. (c) Height recovered using traditional Lambertian photometric stereo. Note the shearing of the surface which is typically due to ignoring non-Lambertian effects. (d) Height recovered using the theory of Section 8.2.

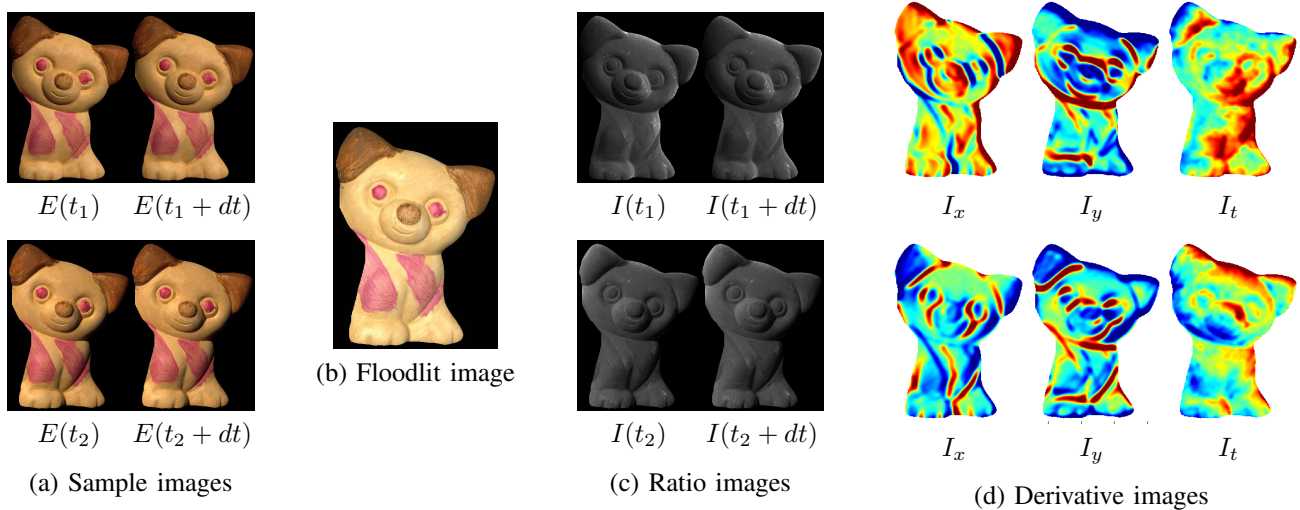


Fig. 15. Input pipeline for the toy dog dataset. (a) Sample images from the differential pair. (b) Floodlit image. (c) Ratios with respect to the floodlit image. (d) Spatial and temporal derivatives. Red stands for positive values and blue stands for negative values.

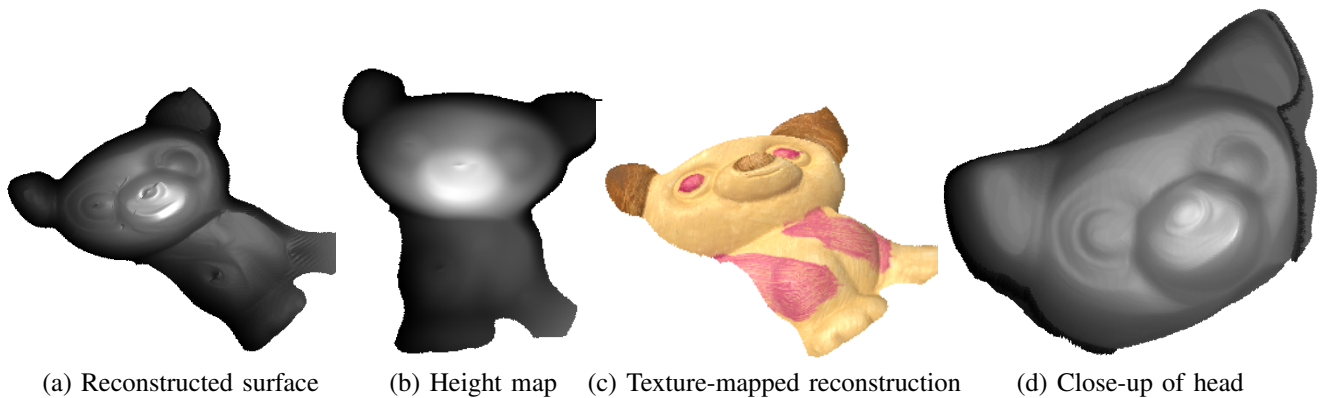


Fig. 16. Reconstruction for the toy dog using Algorithm 2 of Section 8.2, with Dirichlet boundary conditions. Note that the surface is not continuously differentiable, so does not strictly meet the requirements of our theory. Yet, the PDE solution is robust enough to recover a high quality depth map. (a) A top-and-side view of the recovered surface. (b) Recovered height map. Brighter shading indicates higher values. (c) A texture-mapped display of the 3D reconstruction. (d) A close-up reconstruction of the head.

Note the fine scale structure recovered by the algorithm, such as the eyes and the smiling mouth. There is a minor loss of detail near the feet of the toy dog, which might be due to shadows and low levels of illumination (all the light source positions are above the level of the head). While the theory of the paper is derived for differentiable surfaces, this example shows robustness to minor non-differentiability, provided the surface is continuous. We remind the reader that all the reconstructions in this paper are obtained without any calibration for light source positions.

9 IMPLEMENTATION DETAILS

In this section, we provide some implementation details relevant to the real data experiments presented in the previous sections.

9.1 PDE Solution

We take central differences to numerically solve the system of equations arising from (48) at every pixel:

$$\begin{aligned} z_x &= \frac{z_r - z_l}{2\Delta x}, & z_{xx} &= \frac{z_r - 2z + z_l}{(\Delta x)^2}, \\ z_y &= \frac{z_t - z_b}{2\Delta y}, & z_{yy} &= \frac{z_t - 2z + z_b}{(\Delta y)^2}, \end{aligned}$$

where $\{\Delta x, \Delta y\}$ are the step lengths in the x and y directions and $\{z_l, z_r, z_b, z_t\}$ are the pixels on the left, right, bottom and top of z , respectively, which are one step length away. The depths of the boundary pixels are specified to be 0. Thus, we obtain a linear system of the form $\mathbf{A}\mathbf{z} = \mathbf{b}$, where \mathbf{A} is a large, but extremely sparse matrix. Note that the reconstructed depths at the boundaries may be non-zero, since the PDE constraints and the boundary constraints are solved in a single linear system, in a least squares sense.

Points where $|\lambda|$ and $|\kappa|$ cross a certain threshold (set to 50 in our implementation) are detected. We do not write the PDE constraint at these points, instead, we impose a condition that the depth must be continuous at these points.

Optionally, a regularization term (such as an isotropic prior on the norm-squared gradient, or the Laplacian) may be added to enhance smoothness of the solution. The regularized objective function is

$$\min_{\mathbf{z}} \|\mathbf{A}\mathbf{z} - \mathbf{b}\|_2^2 + \mu \|\nabla \mathbf{z}\|_2^2. \quad (49)$$

When regularization is used, we minimize (49) using the CVX convex optimization software [6]. For the apple dataset, we use a value of $\mu = 0.01$. For the teflon ball dataset, no regularization is required ($\mu = 0$).

For the toy dog dataset, instead of the above regularization, we weight each PDE by a multiplicative factor $\frac{1}{\min\{1 + |\lambda|, 1 + |\kappa|, \tau\}}$. This downweights the contribution of regions where λ and κ tend to infinite values. We use a value of $\tau = 5$, although the optimization can tolerate a large range of values.

9.2 Computing Derivatives

An important consideration in derivative-based methods is handling the noise-amplifying nature of differentiation. Thus, we compute spatial image derivatives with a smoothing kernel, such as a Savitzky-Golay filter [15].

Computing temporal derivatives is more challenging, since one cannot take advantage of a wide support region available for the spatial smoothing filters. Instead, we approximate the temporal derivative by individually smoothing the two images of the differential pair and computing the difference.

9.3 Acquisition Setup

The images for real data experiments are acquired using a Canon EOS 5D Mark II camera. For the apple (Figure 12), teflon ball (Figure 14) and toy dog (Figure 15) datasets, we use a gantry to acquire images at Δt corresponding to 2 degrees.

However, the fact that we need only unconstrained motion of the differential pair on a circle suggests a simple acquisition setup that does not require an expensive gantry. The alternative acquisition setup is used to acquire images for the clay ball dataset of Figure 11(a) and is pictured in Figure 17. It consists of two light bulbs mounted close together at the end of a rigid shaft. The length of the shaft is set to 8 times the distance between the two lights, which results in a Δt of $\frac{\pi}{8}$ radians (approximately 7 degrees). Note that an unmarked piece of string suffices to create a Δt corresponding to any integral ratio of π .

The shaft is connected to a wheel, which can be rotated using the crank. A camera is placed facing the object, aligned with the axis of rotation of the wheel. The crank is moved by hand and rigidly clamped to a table while acquiring images, resulting in different unknown positions of the differential light pair. The differential image pairs are acquired by turning lights 1 and 2 on and off at a few different positions of the crank.

10 DISCUSSION AND FUTURE WORK

In this paper, we have presented a comprehensive theory that relates image gradients to surface geometry for unknown, isotropic BRDFs. We have presented novel invariants for surface reconstruction and precisely characterized the extent to which they inform surface recovery in shape from shading and uncalibrated photometric stereo. In the process, we have uncovered fundamental insights into the nature of differential information contained in photometric images for isotropic BRDFs. Indeed, our theory also places fundamental limitations on any inference problem, not just limited to photometric surface reconstruction. Namely, using image derivatives in the presence of unknown BRDF, without any additional priors, is equivalent to deriving information from the ambiguous topological classes we outline.

A key observation in our work is the linearity of the differentiation operation, that may be used to derive novel constraints on surface geometry, regardless of the exact form of the BRDF. This insight is of potential relevance in many other domains like optical flow or multiview stereo and may provide a unified framework to analyze all such problems with



Fig. 17. A simple acquisition setup that exploits the fact that our approach does not require calibrated light source positions.

general BRDFs. Theoretical analysis of such a framework is already the subject of our ongoing work.

ACKNOWLEDGMENTS

This work was funded by ONR PECASE grant N00014-09-1-0741, a National Science Scholarship from the A*STAR Graduate Academy of Singapore and generous support from Intel, Adobe, NVIDIA and Pixar. We gratefully acknowledge the resources provided by Prof. Marc Levoy and help from Andrew Adams towards data acquisition using the Stanford Spherical Gantry. We thank Miloš Hašan for the bunny data, Neil Alldrin for the code from [1], Profs. Todd Zickler, Steve Marschner and Szymon Rusinkiewicz for insightful discussions and the anonymous reviewers of [4] for helpful comments.

REFERENCES

- [1] N. Alldrin and D. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance : A stratified photometric stereo approach. In *IEEE International Conference on Computer Vision*, 2007. 2, 8, 9, 14
- [2] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2008. 2
- [3] S. Barsky and M. Petrou. The 4-source photometric stereo method for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(10):1239–1252, 2003. 2
- [4] M. Chandraker, J. Bai, and R. Ramamoorthi. A theory of differential photometric stereo for unknown isotropic BRDFs. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2505–2512, 2011. 2, 14
- [5] J. Clark. Active photometric stereo. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 29–34, 1992. 3
- [6] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming. <http://cvxr.com/cvx>, 2010. 13
- [7] A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1254–1264, 2005. 2
- [8] M. Holroyd, J. Lawrence, G. Humphreys, and T. Zickler. A photometric approach for estimating normals and tangents. In *SIGGRAPH Asia*, 2008. 2

- [9] B. Horn. *Robot Vision*. MIT Press, 1986. 3
- [10] K. M. Lee and C.-C. J. Kuo. Shape from shading with a generalized reflectance map model. *Computer Vision and Image Understanding*, 67:143–160, 1997. 4
- [11] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Image Understanding Workshop*, pages 121–130, 1981. 3, 6
- [12] M. Brooks and B. Horn. Shape and source from shading. In *International Joint Conferences on Artificial Intelligence*, 1985. 3
- [13] E. Prados and O. Faugeras. Shape from shading: A well-posed problem? In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 870–877, 2005. 4
- [14] R. Ramamoorthi, D. Mahajan, and P. Belhumeur. A first order analysis of lighting, shading and shadows. *Transactions on Graphics*, 26(1), 2007. 3
- [15] A. Savitzky and M. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964. 13
- [16] P. Tan, S. Mallick, D. Kriegman, L. Quan, and T. Zickler. Isotropy, reciprocity, and the gbr ambiguity. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2007. 2
- [17] P. Tan and T. Zickler. A projective framework for radiometric image analysis. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 2, 7
- [18] P. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980. 2
- [19] R. Zhang, P. Tsai, J. Cryer, and M. Shah. Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999. 3
- [20] R. Zhang, P.-S. Tsai, J. E. Cryer, and M. Shah. Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999. 4, 5
- [21] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision*, 49(2/3):1215–1227, 2003. 2